IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of: § Serial No: 09/233,860

Scott H. Hutchinson and §

Gregory M. Hanka § Filed: January 20, 1999

§

For: SOFTWARE-IMPLEMENTED METHOD § Group Art Unit: 2763

FOR IDENTIFYING NODES ON A §

NETWORK §

§ Examiner: William D. Thomson

§ 703.305.0022

Atty Dkt: 23946-P001US §

Assistant Commissioner for Patents §

Washington, DC 20231 §

SECOND SUBSTITUTE APPEAL BRIEF

TABLE OF CONTENTS

TABLE OF AUTHORITIES

I.   **INTRODUCTION**

This is an appeal to the Board of Patent Appeals and Interferences of the rejection of claims 1-3, 5-8, 9-13, and 15-24 in a "final" Office Action dated May 10, 2000 ("the Office Action") in application S.N. 09/233,860 ("the '860 application"). A Notice of Appeal was submitted on August 3, 2000 and received by the Patent & Trademark Office on August 7, 2000, such that the deadline for submission of an Appeal Brief was October 7, 2000. An Appeal Brief ("the Appeal Brief") was filed on October 3, 2000, along with the requisite fee payment pursuant to 37 C.F.R. §§ 1.192 and 1.17(c).

In an Office Action dated January 2, 2000 ("the Office Action"), the Appeal Brief was objected to as failing to comply with various provisions of 37 C.F.R. § 1.192. These objections necessitate the submission of this Substitute Appeal Brief.

A.   **Real Party in Interest**

The real party in interest in this Appeal is BindView Development Corporation, a Texas corporation having a place of business at 5151 San Felipe, Suite 2100, Houston, Texas 77056, by virtue of an assignment of the application dated January 20, 1999 and recorded at Reel 9726, Frame 0990.

B.   **Related Appeals and Interferences**

None.

C.   **Status of Claims**

Claims 1-3, 5-8, 9-13, and 15-24 are pending in the application. Claims 1-3, 5-8, 9-13, and 15-23 have been twice rejected; claim 24, added in an Amendment submitted subsequent to a first Office Action, has been once rejected. Claims 4, 9, and 14 were canceled in a June 11, 1999 "(Preliminary) Amendment A." Claim 17 was canceled in a February 14, 2000 "Amendment B,

Response to November 12, 1999 Office Action." Claims 4, 9, 14, and 17 remain canceled. Claims 1-3, 5-8, 9-13, and 15-24 are appealed.

Pursuant to 37 C.F.R. § 1.192(c)(9), the claims involved in this appeal are reproduced in Appendix A.

D.   Status of Amendments

There are no pending, unentered amendments.

E.   Summary of the Invention

The invention disclosed and claimed in the '860 application is directed to a method and apparatus for asset tracking and management in a computer network. Specifically, the invention relates to "[a] method ... of transmitting asset-management information about [a] node" in a network. Claim 1, lines 1-3.[1] As noted in the specification of the '860 application ("the Specification"), the term "asset management" refers to the process of "track[ing] computers and similar equipment ('nodes'), and their components, on computer networks." (Specification, p. 2, lines 16-18). As will be discussed below, this process is to be distinguished from <u>network</u> management, a term which is known to those of ordinary skill in the art to involve the management of a network's logical topology, constituency, and performance.

The invention disclosed and claimed in the '860 application involves the detection of one or more "unique attribute values" of a client node (i.e., one of the hardware components such as a computer workstation) attached to the network, and the communication of such unique attribute values to a central server program. *See, e.g.,* Specification p. 4, lines 14-16. In accordance with the disclosed invention, "[t]he one or more unique attribute values are also stored to a local database at the client node." (*Id.,* p. 4, lines 18-19).

---

[1] Unless otherwise noted, all citations to claim line numbers herein shall refer to the claims as they appear in the "Amendment B, Response to November 12, 1999 Office Action" submitted February 14, 2000 in connection with the present application. Citations to page and line numbers in the Specification shall refer to the application as originally filed.

In one disclosed embodiment of the invention, a specific attribute value tracked by the asset management software is the current address of the node's network interface card ("NIC"), along with a former NIC address held by the node, if any. (*Id.*, p. 4, lines 26-29). At the same time, it is noted in the Specification that, without the benefit of the teachings of the Specification, NIC addresses are imperfect attributes to be used for the purposes of asset management, owing to their relatively transient nature in the context of most practical applications. As noted in the Specification, "[a] NIC... is often not a permanent part of a microcomputer's motherboard...; very often it is a removable component [such that] any asset management product relying solely on the NIC address for node identification will falter when a node's NIC... changes.... By analogy, the FBI would have a similar problem if a person's fingerprints were to change every time the person got a manicure." (Specification, p. 8, lines 3-13).

A feature of the subject invention noted in the Specification relates to the ability of an asset management system in accordance with the invention to reliably identify the particular client nodes making up or coupled to the network despite the fallacies of nearly all metrics for tracking hardware components. Table 1, appearing on page 10 of the Specification, summarizes the various available metrics and identifies each one's shortcomings. The Specification notes with reference to Table 1 that at present, the least fallible metric for component tracking is a fledgling motherboard serial number standard that has yet to gain widespread industry acceptance. (*See*, Specification, p. 6, lines 7-25).

The disclosed invention, therefore, involves a system whereby the efficacy of tracking a network node's NIC address is substantially augmented through the introduction of a protocol which accounts for unpredictable and otherwise untraceable changes in a component's NIC address. Through application of the teachings of the subject invention, effective asset management can be realized notwithstanding the shortcomings of prior art techniques.

The invention achieves its objectives through the use of "asset-management information" associated with each node on the network. *See, e.g.*, claim 1, line 2; claim 8, lines 4-5; claim 21, line 10; claim 23, line 8; claim 24, lines 2 and 5. In particular, it is to be noted that the

3

present invention contemplates that such "asset-management information" comprises information *in addition to* a node's current and/or former NIC address. Claim 1, for example, recites "transmitting asset-management information concerning the node *together with* the current NIC address value and the former NIC address value." (Claim 1, lines 8-9). It is the **combination** of the "asset-management information" with the current and former NIC address values that enables the invention to uniquely identify nodes on a network; and, as will be discussed hereinbelow, it is this "asset-management information" that is lacking from the prior art of record.

Independent claims 8 and 21 similarly recite "asset-management information" in combination with current and former NIC address values. Independent claim 16 uses the alternative phraseology "node-identification information."

Independent claim 16 further specifies an additional or alternative aspect of the invention, namely that the current and former NIC address values are each associated with a timestamp. *See,* claim 16, line 6. This too is lacking in the prior art.

F.   Issues

(a)   Are claims 11 and 12 indefinite within the provisions of the fourth paragraph of 35 U.S.C. § 112 because they fail to further limit the claims from which they depend?

Assignee answers "no."

(b)   Are claims 11 and 12 indefinite within the provisions of the second paragraph of 35 U.S.C. § 112 for failing to particularly point out and distinctly claim the subject matter which Applicant (Assignee) regards as the invention?

Assignee answers "no."

(c)   Are claims 1-3, 5-8, 9-13, and 15-24 properly rejected under 35 U.S.C. § 102 because they are anticipated by the cited prior art?

Assignee answers "no."

G.   Grouping of Claims

Assignee argues hereinbelow the patentability of the claims under the following groupings according to the patentability issues presented.

35 U.S.C. § 112 ¶ 2: Claims 11-12.

35 U.S.C. § 112 ¶ 4: Claims 11-12.

35 U.S.C. § 102: Claims 1-3, 5-8, 9-13, and 15-24.

As to the § 112 rejections, claims 11 and 12 stand or fall as a group. As to the § 102 rejections, claims 1-3, 5-8, 9-13, and 15-24 stand or fall as a group.

H.   Argument

1.   35 U.S.C. § 112 ¶ 2: Claims 11-12

Claims 11 and 12 stand rejected under 35 U.S.C. § 112, second paragraph. For convenience of reference, exemplary claim 11 is reproduced here:

> 11.   A program storage device readable by a processor in the client node of a specified one of claims 1 through 3, 5 through 7, and 21 through 24, and encoding a program of instructions including instructions for performing the operations recited in the specified claim as being performed by the client node.

(Amendment B, submitted February 14, 2000).

Because of this claim's dependance from claim 1 (*inter alia*), claim 1 is likewise reproduced here for convenience of reference:

> 1.   A method, executed by a node on a network, said node comprising at least one asset, of transmitting asset-management information about the node, the method comprising:
> (a) determining a current address value of a network interface card of the node, referred to as a NIC address value;
> (b) retrieving, from a data storage at the node, a former NIC address value for the node; and
> (c) transmitting asset-management information concerning the node together with the current NIC address value and the former NIC address value.

Id.

According to the Office Action, "[claims 11 and 12] are written in a manner that does not distinguish them as either method or computer readable medium, but rather some type of hybrid wherein the computer readable medium cannot be clearly correlated to specific method steps." This assertion is respectfully traversed.

As a preliminary observation, it is to be noted that the propriety of claim formulations such as adopted by claims 11 and 12 is well established and has much precedent under U.S. patent law. Although the argument has been previously raised earlier in the prosecution of the present application (in an "Amendment B, Response to November 12, 1999 Office Action and Summary of Telephone Interview" submitted February 14, 2000), it bears worth repeating in this forum: the Court of Appeals for the Federal Circuit has approved at least one claim in the same basic format as that of claims 11 and 12 at issue here. In *In re Warmerdam*, 33 F.3d 1354 (Fed. Cir. 1994), the Federal Circuit reversed an indefiniteness rejection, noting that "[t]here is no requirement that a claim for a machine which incorporates process steps... must conform to the conventional definition of a product-by-process claim," holding that "[t]here has been no showing that one skilled in the art would have any particular difficulty in determining" the scope of the claim in question. *Id.*

With regard to the claims at issue in this appeal, there is a similar lack of indefiniteness. Claim 11 refers to a "program storage device," for example, a computer disk, and specifies only that the device "encod[es] a program of instructions for performing the operations recited" in certain method claims. It is respectfully submitted that such a claim formulation gives rise to no indefiniteness. Claims 11 and 12 are arguably nothing more than multiple-dependent claims;[2] the Office Action itself notes that "the applicant has been assessed a surcharge for this type of multiple dependent claim language." (Office Action, p. 5).

---

[2]   It is not clear to the Assignee that claims 11 and 12 are multiple-dependent claims. Arguably, claims 11 and 12 are Markush-type claims. However, the specific label applied to claims 11 and 12 is immaterial to the substance of the arguments advanced herein.

Consideration of an isolated sub-part of claim 11 is illustrative of the lack of indefiniteness: As applied to claim 1, claim 11 recites as follows:

> "A program storage device readable by a processor in the client node of ... claim 1... encoding a program of instructions including instructions for performing the [following] operations: (a) determining a current ... NIC address value; (b) retrieving, from a data storage at the node, a former NIC address value for the node; and (c) transmitting the current NIC address value and the former NIC address value."

Pared to its essentials, it is entirely clear what is claimed:"A program storage device readable by a processor in the client node of claim 1" (a simple dependent claim) "encoding a program of instructions... for performing" specified operations. The multiple-dependent nature of claim 11 gives rise only to other equally definite claims, a privilege for which the Assignee has already paid. It is thus submitted that nothing about the resultant claim language is in any way indefinite. That claim 11 claims as alternatives the incorporation of other method claims renders claim 11 nothing more or less than a multiple dependent claim, the propriety of which being well established. Hence, there is no justification for the Office Action's assertions that "claim 11 cannot be examined, allowed, or rejected in total." Examination of claims 11 and 12 is no more "difficult if not impossible," as the Office Action alleges (Office Action, pp. 5 & 6), than any other multiple-dependent claim.

The Assignee further notes that the formulation of claims 11 and 12 is undeniably precedented. A search of the U.S. Patent & Trademark Office's Internet web site (www.uspto.gov) has uncovered more than 500 issued patents having comparable claim formulations.[3]

---

[3]   Using the Boolean search terms "program storage device" and "tangibly embodying" in the claim language of all U.S. patents issued since 1976, more than 500 patents were identified. A listing of the search results is attached as Exhibit B. Among those patents listed in Exhibit B is U.S. Patent No. 5,860,929 to Rubin et al., which includes a claim exemplary of those at issue here; a copy of the '929 patent is attached as Exhibit C. Claim 11 of the '929 patent recites in its entirety: "[a] program storage device readable by a processor in an ultrasound machine, tangibly embodying a program of instructions executable by the processor to perform the method of a specified one of claims 1 through 3." It is submitted that such language is indistinguishable from that of the claims at issue herein.

2. 35 U.S.C. § 112 ¶ 4: Claims 11-12

The Office Action's assertion that claims 11 and 12 are invalid under paragraph 4 of 35 U.S.C. § 112 is simply not understood and, accordingly, respectfully traversed. Claims 11 and 12 are, distilled to their substance, directed to "program storage devices," whereas the respective claims from which they depend are directed to "a method ... of transmitting [or recording] asset-management information concerning [a network node]." Given that the claims from which claims 11 and 12 depend recite methods, whereas claims 11 and 12 recite "program storage device[s]" (e.g., hard disk drives, diskettes, and the like), the Office Action's assertion that claims 11 and 12 "imply the same scope relative to the claims to which each depends" finds no basis in fact.

As has been noted in prior submissions relating to the present application, claims 11 and 12 are designed to read upon the physical media upon which certain program instructions may be stored, as distinctly contrasted with the method implemented by such instructions. It is inarguably improper to assert that there is identity between the physical embodiment (a computer disk, e.g.,) and an intangible methodology. Not only do claims 11 and 12 "further limit" the claims from which they depend, they wholly transform those claims into new scopes of coverage.

3. 35 U.S.C. § 102: Claims 1-3, 5-8, 10-13, 15-16, and 18-24

In the Office Action, claims 1-3, 5-8, 10-13, 15-16, and 18-24 were rejected under 35 U.S.C. § 102 as being unpatentable over U.S. Patent No. 5,878,420 to de la Salle ("*de la Salle*") and U.S. Patent No. 5,923,850 to Barroux ("*Barroux*").[4]

The patentable distinctions between the invention disclosed and claimed in the present application from *de la Salle* and *Barroux* shall be separately discussed below. There is at least one feature common to both references, however, that merits preliminary notice. In particular,

---

[4] The *de la Salle* and *Barroux* references were not cited in the first Office Action (dated November 12, 1999) in connection with the subject application. Claim 24, added only after issuance of the first Office Action, falls clearly within the scope of the disclosure and claims of the application as originally filed. Hence, it is not clear to the Assignee the basis for the Office Action's statement that "Applicant's amendment necessitated the new ground(s) of rejection presented...." Likewise, the basis for the designation of "finality" of the (second) Office Action is not understood.

8

the Assignee submits that both of the cited references assume away the very essence of the problem to which the present invention is directed.

As noted above, the present invention is directed to an asset management system for a computer network. A key feature of the invention is the ability of the system to uniquely identify each node on the network. As noted in the specification, "[t]here does not yet [at least at the time of filing the '860 application] exist a standard, ubiquitous 'fingerprint' for computers... so asset management products must approximate one using whatever shifting data they can find on each node." (Specification, p. 6, lines 3-5). The Specification discusses in detail various attributes of network nodes which might be considered useful for the purposes of uniquely identifying them, and as summarized in Table 1 on page 10, concludes that only one attribute -- one that has not found widespread acceptance among computer equipment manufacturers -- was not susceptible to failure as a node identifier.

In the cited references, on the other hand, the existence of unique identifiers for network nodes is not discussed; rather, it is assumed. That is, as shall be discussed in greater detail below, the cited references completely ignore the problem sought to be solved by the present invention, and indeed take as their respective fundamental premises that such a problem does not exist. As a result, each cited reference fails to teach or suggest critical elements of the claimed invention.

(a) U.S. Patent No. 5,878,420 to de la Salle

The very title of *de la Salle* -- "Network Monitoring and Management System" (emphasis added) -- reveals that this reference is not directed to the problem sought to be solved by the present invention, namely asset monitoring and management. *de la Salle* asserts that "a primary task [of a network management system] is to keep track of the actual configuration of the network and, following that, to reconfigure or otherwise optimize or 'tune' the network, as necessary, so as to minimize problems and maximize the utilization of resources." *de la Salle*, col. 1, lines 58-63. That is to say, *de la Salle* is concerned only with the configuration and topology of a network

comprising a plurality of interconnected components, rather than with the identity of specific hardware components making up the network. While *de la Salle* seeks to obtain and track information useful for the purposes of enhancing and maximizing overall network performance, the present invention seeks to obtain and track information useful for the purposes of asset management.

Because *de la Salle*'s objectives are entirely different than those promoted through the practice of the present invention, *de la Salle* is unconcerned with the very problem sought to be solved by the present invention. It is clear that *de la Salle* regards a network component's NIC address (*de la Salle* uses the term "board address") as a sufficiently unique identifier. That is, *de la Salle* fails to teach or suggest "asset-management information" as disclosed and claimed in the present application. de *la Salle* observes that "[o]rdinarily, each network component 16 will have a single network address which is used by the system in order to locate that particular component 16." *de la Salle*, col. 5, lines 33-36. Further, *de al Salle* notes that "[t]he board address 36 is the typical information which is utilized within the network to identify and locate (logically) the particular network component 16." *de la Salle*, col. 5, lines 46-48.

Notably, the present invention specifically identifies the disadvantages of using merely a network address as a node identifier for the purposes of asset management: "The NIC 107, however, is often not a permanent part of a microcomputer's motherboard 108; very often it is a removable component plugged into the motherboard.... [T]hus any asset management product relying solely on the NIC address for node identification will falter when a node's NIC 107 changes in this way." Specification, p. 8, lines 3-11.

de la Salle's exclusive reliance upon each network components "board address" as a unique identifier is further evident from its proposed treatment of network components having no board address (such as bridges), or multiple board addresses (such as routers). *de la Salle*'s discussion of these special cases is highly revealing of *de la Salle*'s exclusive interest in mapping network topology and configuration, as contrasted with the present invention's concern with identification of specific hardware assets.

10

In the former case, *de la Salle* notes as follows with respect to network components having no network address:

> The next process step is in the nature of a resolve bridged stations conflict step 122. This occurs when stations (identical board addresses 36) are detected on different branches of the network 16. When this occurs, it may be *assumed* that the branches 14 are connected by a bridge 28, which (as discussed above) is transparent to the packets 22, since bridges 28 are not assigned network addresses 23 and will not have individual board addresses 36. A heuristic analysis is utilized to '*guess*' the identity of a branch 14 upon which the station is actually located, and a location is assigned as a result of this guess."

de la Salle, col. 11, lines 34-44 (emphasis added)

Again, it is clear that *de la Salle* is wholly reliant upon "board addresses" to be as close to a "unique identifier" of a network component as any available characteristic or parameter. Nonetheless, *de la Salle* must resort to "guessing" in the event that a board address proves unreliable for this purpose -- an unreliability that the present application notes to be inherent. Indeed, were a truly unique identifier such as is obtainable through the practice of the present invention contemplated by and available to *de la Salle*, no such "guessing" would be necessary.

With regard to the latter case of network components having more than one "board address," *de la Salle* states as follows:

> Since routers 30 will have different addresses (both network and board) for each branch 14 to which they are connected, a true picture of the overall network must include some method of resolving these multiple address locations into a single station. For this purpose, a consolidate routers: Phase 1 step 124 and a consolidate routers: Phase 2 step 126 are included in the db builder routine 96. The phase 1 step 124 invokes reasoning based on the standard router query under the SNMP (Simple Network Management Protocol). This involves address analysis in which a single location appears to 'own' and advertise overlapping addresses on different branches. When this occurs, it is *a relatively safe assumption* that the station is a single station component 16 and that it is a router 30. The phase 2 step 126 includes hop count analysis to *postulate* router identity."

de la Salle, col. 11, lines 45-59 (emphasis added).

11

Yet again, it is evident that *de la Salle* regards board addresses as the most reliable distinguishing characteristic of any given network component, and yet again, *de la Salle* proposes nothing more effective than "assumption" and "postulation" to overcome the fallacies of board addresses serving as unique identifies. Applying the teachings of the present application, on the other hand, no such imprecise methods would be necessary.

A further critical deficiency of *de la Salle* relates to its failure to teach or suggest the concept of associating time or date stamps with information obtained from network nodes. It is respectfully submitted that introducing this additional temporal dimension to a node identifier is one of the features of the present invention which sets the invention patentably apart from the prior art, including *de la Salle*. Nowhere does *de la Salle* teach or suggest associating time or date data with the network analysis data it describes. On the contrary, *de la Salle* seems to contemplate completely eliminating the temporal dimension from its network data, and calls for a system in which "the sampling assembly will be continuously providing new probe objects 52 and the analysis assembly will continually enhance and update the database 99, the database management 138 will continually provide fresh and current information on the precise state of the network 12." *de la Salle*, col. 14, lines 42-46.

In fact, *de la Salle* appears to <u>require</u> continuous operation in order to effectively accomplish its objective of monitoring a network's configuration and topology. If the *de la Salle* system were deactivated for any appreciable amount of time (i.e., sufficient for any changes in network topology or configuration to take place), the *de la Salle* system would have to in essence be re-initialized (not *de la Salle*'s terminology), in which case changes in hardware such as are detectable using the system of the present invention would be overlooked. *de la Salle* notes that "[i]n initial operation, the system 10 will require a period of time to sample sufficient data in order to build a working database." *de la Salle*, col. 14, lines 29-31. Thereafter, *de la Salle* contemplates the continuous operation noted above. Because *de la Salle* eliminates any temporal component from the data it collects, its proposed system cannot accomplish either its own contemplated objectives (network monitoring) or those contemplated by the present invention (asset monitoring) unless it

operates continuously. No such continuous operation is necessary in accordance with the teachings of the present application.

A practical example is perhaps the most straightforward way in which to highlight the clear distinctions between *de la Salle* and the present invention. As is known, a network interface card ("NIC," to use the language of the present application, "network board" to use the language of *de la Salle*) can be readily installed in many different computers. The system of the present invention is specifically adapted to detect situations in which a NIC is removed from one computer and installed in a second, even if the second computer is reattached to the network at the same location as the first. The *de la Salle* system, on the other hand would be incapable of detecting such a hardware swap. That is, the particular hardware associated with any given node of a network is of no concern to the *de la Salle* system, only the existence of the node and the presence of some hardware -- any hardware -- at that node. In stark contrast, the identity of particular hardware is precisely the information the present invention seeks to track.

Considering the specific claim language, the deficiencies of *de la Salle* are undeniably apparent. As a starting point, it is to be noted that independent claims 1, 8, 19, 23, and 24 each explicitly recite "asset management information." *de la Salle*, by its own terms, does not disclose an asset management system, but a network management system. Those of ordinary skill in the art will readily appreciate that this is far from a semantic distinction, given that the objectives of one are distinct and unrelated from those of the other. A network management system is concerned such issues as with "determining the configuration of an expansive network" (*de la Salle*, col. 3, lines 11-12), "ascertaining a network configuration and functionality" (*id.*, lines 14-15), "determining the functional and performance characteristics of a computer network" (*id*, lines 19-20), and "provid[ing] a monitoring and management system which compiles network parameter information" (*id.*, lines 21-23). A network management system is relatively unconcerned with the particulars of the hardware present at each node, but rather with optimizing the operation of whatever hardware is present. An asset management system, on the other hand, is concerned with "what equipment is connected to a network," since such information "is necessary for user support, network

13

planning, corporate financial statements, software purchasing" and the like. *Balloux*, col. 1, lines 21-24.

Each of the independent claims 1, 8, 13, 16, 19, 21, 23, and 24 calls for node identification or asset-management information comprising a "current NIC address value" and a "former NIC address value." Simply stated, nowhere does *de la Salle* teach or suggest retrieving a former value from a network node. Hence, *de la Salle* cannot be characterized as anticipating the independent claims. That is, even if the critical distinction between "asset management" and "network management" is (improperly) ignored, each of the claims recites elements that are neither taught nor suggested by *de la Salle*.

In summary, whereas *de la Salle* is concerned with the question "How is the network configured, and which component of a general class of functionality is connected and communicating with which," the system in accordance with the claimed invention asks the question: "Which particular components at a given location are present, and how might their physical configuration have changed since the last time their status was established?" The answer to the former question would be invisible to systems in accordance with the *de la Salle*, whereas the answer to the latter question is the fundamental objective of the subject invention. In view of these critical distinctions, the claim rejections based upon *de la Salle* cannot stand.

(b)   U.S. Patent No. 5,923,850 to Barroux

Unlike *de la Salle*, *Barroux* does purport to be an asset management system rather than a network management system. And unlike *de la Salle*, *Barroux* does appear to associate some type of time or date stamp information with the network node information collected. It might therefore be tempting upon casual consideration to regard *Barroux* as anticipating or rendering obvious the invention disclosed and claimed in the present application, as did the Office Action. However, as shall be discussed below, *Barroux* suffers from deficiencies no less critical -- and perhaps more so -- than *de la Salle*, and provides no support for rejection of the claims at issue.

Barroux appears to disclose a system adapted to monitor both hardware and software components of a network. The *Barroux* system is based upon an "integrated resource 200 for collecting and managing survey information about a network 202...." *Barroux*, col. 3, lines 24-25. The integrated resource "takes advantage of various TCP/IP services and remote execution of commonly installed procedures to automatically learn about nodes of network 202 [and] collects and analyzes information about [network nodes] and returns that information to [an] asset database 232." *Barroux*, col. 3, lines 43-46, and col. 4, lines 10-12.

As can be observed from Figures 7A, 7B, 7C, and 7D, the asset database contemplated by *Barroux* includes information about the "systems" (Figure 7A), "processors" (Figure 7B), "software packages" (Figure 7C), and "patches" (Figure 7D) present on the network under analysis. *Barroux* further suggests that information is similarly gathered for memory, buses, peripherals, and interfaces in the system. *See Barroux*, col. 10, lines 9-10. As can further be observed from the referenced Figures, each table includes multiple fields of information. Careful consideration of the information maintained in the various *Barroux* information tables exposes a fundamental deficiency of *Barroux*, namely that, like *de la Salle, Barroux* assumes away the very problem sought to be solved by the present invention.

With reference to Figure 7A, *Barroux* describes beginning at col. 8, line 65 through col. 9, line 32, the various fields in the "systems table" storing information about "host systems" (computers) on the network. A first field, designated AALID, identifies an IP address of a node on the network. (As an aside, it is well-known to those of ordinary skill in the art that IP addresses may be assigned to networked computers on a highly dynamic basis, making an IP address wholly unsuited to serve as a "unique identifier" of a piece of computer hardware.) Of particular significance, however, is *Barroux's* reference to "a CID number which *uniquely identifies* the system operating at the referenced node" *Barroux*, col. 9, lines 7-8 (emphasis added). With this modest statement, *Barroux* utterly assumes away the entire issue about which the present application is concerned. A substantial portion of the Specification of the present application, spanning from page 5, line 27 through page 13, line 9, is devoted to a discussion of the practical unavailability in present

15

systems of the unique node identifier that *Barroux* quite casually and without explanation, assumes to exist!

Similarly, with regard to the "processors table" described with reference to Figure 7B, *Barroux* states, also without explanation, that "each record includes CID field 704" and "a PathTo field 722 that holds a ***unique identifier of the processor on the system*** in standard UNIX device path name format." *Barroux*, col. 9, lines 35-39.

Likewise, with respect to the other tables of information, *Barroux* states as follows:

"For the memory table, PathTo field 722 holds a ***unique identifier*** of the memory units on the host system in path name format." *Barroux*, col. 10, lines 12-14 (emphasis added).

"The peripherals table includes information about individual peripherals on host systems of network 202. PathTo field 722 includes a ***unique identifier*** of the peripheral on the system in device name path format." *Barroux*, col. 10, lines 24-27 (emphasis added).

"The interfaces table includes information about the interface devices on host systems of network 202. PathTo filed 722 holds the ***unique identifier*** of the interface device of the host system in device path name format." *Barroux*, col. 10, lines 32-35 (emphasis added).

A careful study of *Barroux* has not revealed to the undersigned any further description or explanation of the nature of the purported "unique identifiers" seeming to abound in the *Barroux* system.

Of equal significance is *Barroux*'s failure to teach or suggest the introduction of a temporal dimension to the data for the purposes of establishing a unique identifier for nodes connected to a network in the manner taught by the present application. Although *Barroux* does appear to suggest the maintenance of various time and/or date fields among the information collected, this date data is not used as "identifying information" as disclosed and claimed in the present application.

16

According to *Barroux*, several pieces of date data are maintained in the "information tables," including, with reference to the "system table" of Figure 7A, an "FUD field 710 [for holding] a time when the configuration information in configuration field 708 was first observed for the system[, and an] LUD field 712 [for holding] the last time the configuration information in configuration field 708 was observed to be valid." *Barroux*, col. 9, lines 21-25. This date information, however, is used by *Barroux* in order to maintain a configuration "history" for network nodes. *Barroux* states that "[a] change in the contents of configuration field 708 for a given node triggers creation of a new record with a new version number" and that "by examining a series of records for a particular node within system table 700, one can track the system configuration history for each node." *Barroux*, col. 9, lines 29-31 and col. 11, lines 10-12. Nothing in *Barroux* suggests that the date data is used to augment another node-identifier value in order to achieve "unique identification" of a network node. To the contrary, as discussed at length above, *Barroux* simply assumes without discussion that a unique identifier is available for each node.

Thus, while it could be tempting to draw a comparison between the time/date stamping proposed by *Barroux* and the temporal dimension essential to the practice of the present invention, the context in which the *Barroux* time stamping is performed must be carefully considered. Because *Barroux* uses date data merely to provide a configuration history, *Barroux* cannot be said to anticipate the claims of the present application, which call for the transmission (claims 1, 21 and 23), receiving (claim 8) or storage (claims 13, 16, 19 and 24) of "asset management information" comprising a current and former NIC address or node-identifier value. Hence the rejection of the claims based on *Barroux*, like the rejection based on *de la Salle*, cannot stand.

4.   Internet Engineering Task Force Requests for Comments

The Office Action's citation to various Requests for Comment ("RFCs") published by the Internet Engineeering Task Force ("IETF") is acknowledged. It is respectfully submitted,

however, that any inference that the cited RFCs anticipate or render obvious the claims at issue is misplaced.[5]

The cited RFCs appear to relate to the Internet standard "Simple Network Management Protocol," or "SNMP" and related Internet constructs, including CMIP and CMOP. (*See* Office Action, p. 7). As such, the cited RFCs provide no additional basis for the assertions of the Office Action that the subject invention has been anticipated in the prior art. Indeed, the very prior art relied upon by the Office Action (namely, *Barroux*,) states as a premise that SNMP's "primary application has been monitoring network performance rather than asset surveying." *Barrroux*, col. 1, lines 44-47. Asset surveying, on the other hand, is precisely what is disclosed and claimed in the '860 application.

II.   CONCLUSION

In view of the foregoing, it is submitted that the claims in the present application recite combinations of elements neither taught nor suggested by the prior art, including the art relied upon for the purposes of rejection in the Office Action. The present invention is directed to solving a problem that is not even recognized by the cited references, and hence involves structures or method steps not present in the prior art.

de la Salle is directed to a <u>network</u> management system, as opposed to an <u>asset</u> management system. *de la Salle* wholly fails to "identify" network nodes as disclosed and claimed in the present application. Further, *de la Salle* fails to associate any date information with the node information it does collect, such that *de la Salle* does not -- and cannot -- utilize "current" and "former" node-identifying values as required by each independent claim in the present application.

Barroux explicitly assumes away the very problem sought to be solved by the present application, namely the lack of reliable "unique indicators" of network hardware. Although Barroux appears to record date data relating to node information it collects, this is done for reasons wholly unrelated to the problem of node identification. *Barroux* therefore does not transmit (or store)

---

[5]   The RFCs are cited in the Office Action under the caption "Specific Response to Interpretation of Applicant's invention. The Office Action does not rely on the cited RFCs in support of any claim rejection.

"current" and "former" node-identifier values as required by each independent claim in the present application. Thus, *Barroux* cannot be said to teach or suggest the transmission of "asset-management information" as that term must be construed in light of the Specification of the '860 application.

Thus neither *de la Salle* nor *Barroux*, even if considered in a hypothetical combination,[6] teaches or suggests a method or apparatus as disclosed and claimed in the present application.

* * * * *

Reconsideration and withdrawal of the rejections of the claims is therefore requested, such that the application may advance to issue.

Respectfully submitted,

Date: **16-MAY-2001**

Hugh R. Kress
Reg. No. 36,574
Winstead Sechrest & Minick P.C.
2400 Bank One Center
910 Travis Street
Houston, Texas 77002
(713) 650-2714 (voice)
(713) 650-2400 (fax)
ATTORNEY FOR ASSIGNEE

---

[6] The Office Action does not propose a hypothetical combination of the cited references, and no § 103 rejection is advanced.

APPENDIX A

1.   A method, executed by a node on a network, said node comprising at least one asset, of transmitting asset-management information about the node, the method comprising:

      (a) determining a current address value of a network interface card of the node, referred to as a NIC address value;

      (b) retrieving, from a data storage at the node, a former NIC address value for the node; and

      (c) transmitting asset-management information concerning the node together with the current NIC address value and the former NIC address value.

2.   The method of claim 2, wherein determining the current NIC address value includes an attempt to detect the then-current NIC address value.

3.   The method of claim 2, wherein the attempt to detect the then-current NIC address value is unsuccessful, and further comprising (i) retrieving, from a data storage at the node, a stored value containing the result of the past live detection of the then-current NIC address value, referred to as a previously-detected NIC address value; and (ii) transmitting the previously-detected NIC address value.

4.   (previously canceled)

5.   The method of claim 1, wherein the NIC address value comprises a signature portion and a pseudorandomly generated portion.

6.   The method of claim 1, wherein the former NIC address value is redundantly stored in multiple partitions within the data storage at the node.

21

7.   The method of claim 6, wherein (x) each copy of the former NIC address value is associated with a timestamp, and (y) retrieving the former NIC address value comprises retrieving the respective copy associated with the most recent timestamp.

8.   A method, executed by a server node on a network, for recording, in a database, asset-management information about a client node, comprising:

        (a) retrieving, from the client node, (1) asset-management information about the client node, (2) a current address value of a network interface card of the client node, referred to as a current NIC address value and (3) a former NIC address value for the client node that is equal to the current NIC address value;

        (b) unsuccessfully attempting to locate, in the database, a record corresponding to the current NIC address value;

        (c) unsuccessfully attempting to locate, in the database, a record corresponding to the former NIC address value; and

        (d) storing the asset-management information, the current NIC address value, and the former NIC address value in a record in the database associated with the current NIC address value.

9.   (previously canceled)

10.   The method of claim 8, wherein the NIC address value comprises a signature portion and a pseudorandomly generated portion.

11.   A program storage device readable by a processor in the client node of a specified one of claims 1 through 3, 5 through 7, and 21 through 24, and encoding a program of instructions including instructions for performing the operations recited in the specified claim as being performed by the client node.

12.   A program storage device readable by a processor in the server node of a specified one of claims 8, 10, and 24 and encoding a program of instructions including instructions for performing the operations recited in said specified claim as being performed by the client node.

13.   In a node on a network, a data store comprising a machine-readable data structure accessible to a processor in the node and containing node-identification information for the client node that includes (i) a current network interface card value for the node, referred to as a NIC address value, and (ii) a former NIC address value.

14.   (previously canceled)

15.   The data store of claim 13, wherein the NIC address value that constitutes the current node-identifier value includes a signature portion and a pseudorandomly generated portion.

16.   In a node on a network, a data store comprising:
        (a) a plurality of machine-readable data structures accessible to a processor in the node;
        (b) each said data structure containing node-identification information for the client node that includes (i) a current node-identifier value, and (ii) a former node-identifier value, each said value comprising a network interface card address value, referred to as a NIC address value;
        (c) each said data structure being associated with a timestamp.

17.   (previously canceled)

18.   The data store of claim 16, wherein the NIC address value comprises a signature portion and a pseudorandomly generated portion.

23

19.   In a server node on a network, that includes a client node, a machine-readable data structure accessible to a processor in the server node, comprising (i) a current value of a network interface card address for the client node, referred to as a current NIC address value for the client node, (ii) a former NIC address value for the client node, and (iii) asset-management information about the client node.

20.   The machine-readable data structure of claim 19, wherein the current NIC address value comprises a signature portion and a pseudorandomly generated portion.

21.   A method, executed by a node on a network, of transmitting asset-management information about the node, the method comprising:

    (a) determining a current node identifier value, where (1) the node identifier value for any particular node in the network is dependent upon one or more node-identification attributes of that node including an address value of a network interface card in the node, referred to as a NIC address value, and (2) determining the current node identifier value includes an attempt to detect the then-current values of said one or more node-identification attributes;

    (b) retrieving, from a data storage at the node, a former node identifier value for the node; and

    (c) transmitting asset-management information about the node together with the current node-identifier value and the former node identifier value.

22.   The method of claim 21, wherein the attempt to detect said one or more node-identification attributes fails to detect at least one of said node-identification attributes, and further comprising (i) retrieving, from a data storage at the node, a stored value containing the result of a past live detection of the said one or more node-identification attributes, referred to as a previously-detected node identifier value; and (ii) transmitting the previously-detected node identifier value.

23.   A method, executed by a node on a network, of transmitting asset-management information about the node, the method comprising:

(a) attempting but failing to detect a current network interface card address value for the node, referred to as a current NIC address value;

(b) retrieving, from a data storage at the node, a previously-detected NIC address value;

(c) retrieving, from a data storage at the node, a stored value of a former NIC address value for that node; and

(d) transmitting the asset-management information together with the previously-detected NIC address value and the former NIC address value.

24.   A method, executed by a client node and a server node on a network, for recording, in a database, asset-management information about the client node, comprising:

(a) the client node (1) determining a current address value of a network interface card in the node, referred to as a NIC address value, (2) retrieving, from a data storage at the node, a former NIC address value for the node, and (3) transmitting to the server node asset-management information, the current NIC address value, and the former NIC address value;

(b) the server node (1) unsuccessfully attempting to locate, in the database, a record corresponding to the current NIC address value, (2) locating, in the database, a record corresponding to the former NIC address value, (3) recording the asset-management information in said record, and (4) updating the record to correspond to the current NIC address value instead of the former NIC address value.

25